

Project 2

Intro Lecture

Sam McCoy

Cornell CS 4411, September 23, 2016

Today's Lecture

- Administrative Information
- Project 2 Rundown
- Discussion

Administrative Information

- Project 1 is still being graded – results should be ready by beginning of next week.
- Project 2 will be due at 11:59pm on October 6th.

P2 Overview

- Three main parts of P2:
 - Clock Interrupts
 - Add Alarms
 - Implement `thread sleep_with_timeout()`
 - Multilevel Feedback Queue Scheduler

Where are we now?

- We made a very large assumption in P1
 - Threads will behave and yield the CPU to other threads.
 - Issues ?
 - Selfish Threads
 - Starvation
- Need the ability to interrupt running threads...

Introducing Interrupts

- Every clock tick the interrupt handler will be called.
- Allows us to schedule a new thread to be run without relying on nice behavior.
- Maskable Interrupts:
 - `Minithread_clock_init()`
 - `Set_interrupt_level()`
- Always save old interrupt level!

Disabling Interrupts

- When you need to do something that must be done atomically.
- Typically manipulations on shared data structures.
 - Data structures that can be accessed by multiple threads ‘simultaneously’.
 - Modifying the cleanup queue, ready queue, alarm list.
- Trivial way of achieving correctness: disable interrupts for everything.
 - Why is this a bad idea?

Interrupt Handler - Reminder

- Entry point when a clock interrupt occurs.
- Are there problems if the interrupt handler is interrupted?
 - Yes – accessing shared data structures
 - Solution – disable interrupt in the interrupt handler

CANNOT BLOCK

Semaphore Revisited

- Typical sem_P code:

```
while (TAS(&lock) == 1);

sem->counter--;
if (sem->counter < 0)
{
    append thread to blocked queue
    atomically unlock and stop
}
else
{
    atomic_clear(&lock);
}
```

Semaphore Revisited

- Typical sem_V code:

```
while (TAS(&lock) == 1);

sem->counter++;
if (sem->counter <= 0)
{
    take one thread from blocked queue
    start the thread
}

atomic_clear(&lock);
```

Semaphore in User Space

- Interrupts can arrive at any time.
- If interrupts arrive while a TAS lock is held:
 - Another thread that tries to acquire the TAS lock will spin until its time quanta is exhausted.
 - Thread holding the TAS lock will eventually regain control and make progress.
 - Progress ensures other threads can eventually get the TAS lock.

semaphore_P

```
while (TAS(&lock) == 1);

sem->counter--;
if (sem->counter < 0)
{
    append thread to blocked queue
    atomically unlock and stop
}
else
    atomic_clear(&lock);
```

semaphore_V

```
while (TAS(&lock) == 1);

sem->counter++;
if (sem->counter <= 0)
{
    take one thread from blocked queue
    start the thread
}

atomic_clear(&lock);
```

Semaphore In Kernel Space


- Typically used to block some thread and wake it up on some condition
 - `minithread_sleep_with_timeout()`
 - wake up the thread after the elapsed time
- Waking up requires calling `sem_V` on that sleep semaphore
- Where is this done?
 - Done in kernel space with interrupts disabled.

Unfortunate Interleaving


- What if user calls `sleep_with_timeout(0)` ?
 - `sem_P` is called, and thread blocks itself.
- What if `sem_P` was interrupted just after placing thread on blocked queue but before clearing TAS lock?

user calls `sleep_with_timeout(0)`...

```
while (TAS(&lock) == 1);

sem->counter--;
if (sem->counter < 0)
{
    append thread to blocked queue
     clock interrupt!
    atomically unlock and stop
}
else
    atomic_clear(&lock);
```

...clock handler tries to wake that thread up

```
while (TAS(&lock) == 1); 
sem->counter++;
if (sem->counter > 0)
{
    take one thread from blocked queue
    start the thread
}

atomic_clear(&lock);
```

but interrupts are disabled in the clock handler!

Solution

- Disable interrupts for `sem_P` and `sem_V` for `minithread_sleep`
 - Atomicity: `sem_P` will be done with everything before an interrupt can possibly arrive.
 - If interrupt arrives, acquisition of TAS lock is guaranteed in kernel space.
- What about `sem_V`?
 - `sem_V` is called from interrupt handler.
 - Interrupts are already disabled in the handler.

When is this applicable?

- If semaphore will be used in portions of your kernel where interrupts are disabled.
 - Right now: only the sleep semaphore.
- What about cleanup semaphore?
 - Cleanup semaphore is not signaled from any place where interrupts are disabled.
 - Cleanup code should only disable interrupts while accessing the cleanup queue, not for semaphore signaling.

Scheduling




Schedule 80 quanta – 1 quanta per thread



Schedule 40 quanta – 2 quanta per thread



Schedule 24 quanta – 4 quanta per thread



Schedule 16 quanta – 8 quanta per thread

What is the maximum number of unique threads that can run per level?

Scheduling

- Completes 1 sweep over the queue in approximately 160 ticks
- If there are no threads in a given level, schedule threads from the next available level
- Thread level starts at 0 and can only increase throughout its lifetime

Priority Changing

- Threads are scheduled to run for the max duration of the current level
 - For example, in level 1, each thread will be scheduled to run for 2 quanta
- A thread is demoted if it uses up the entire quanta
 - What if the thread is from a different level?

Alarms

- Useful construct for scheduling a thread for future execution
 - Can be used for `minithread_sleep()`
- Each alarms requires a call back function
 - Call back functions might not be executed by the thread that registered the alarm!
- How to keep track of alarms?
 - Add functionality to existing queue.
 - Insert should be $O(n)$, remove min should be $O(1)$.

Alarm Firing

- Where should the alarm be fired?
 - Interrupt handler
- When should an alarm be fired?
 - Tick == alarm expiration time
 - Can this be missed?
- How should be alarm be fired?
 - Context switch to alarm thread?
 - Should fire in the context of the currently executing thread

Testing

- There are a lot of parts to this project
 - Multi-level queue
 - Interrupts
 - Alarms
 - Thread levels
- Common pitfalls
 - Unnecessarily disabling interrupts
 - Not disabling interrupts when necessary
 - Multi-level queue corner cases

Questions

Questions?

Project 2 FAQ

- All library calls are safe: interrupts are automatically disabled upon calling
 - interrupts will be restored to its original state (enabled/disabled) after the call.
- Units of time
 - PERIOD is defined as 50 ms, which is 50000 as a constant.
 - Alarm and wakeup delays are specified in milliseconds.
 - You have to convert units; don't blindly subtract PERIOD.
- Irregular/random clock interrupts
 - This is normal
 - Be careful of introducing heisenbugs because of your debug statements.